arXiv:1704.00565v1 [cs.DS] 3 Apr 2017

Dynamic Planar Embeddings of Dynamic Graphs

Jacob Holm & Eva Rotenberg
Department of Computer Science, University of Copenhagen, Denmark
{jaho,roden}@di.ku.dk

April 4, 2017

Abstract

We present an algorithm to support the dynamic embedding in the plane of a dynamic graph. An edge can be inserted across a face between two vertices on the face boundary (we call such a vertex pair linkable), and edges can be deleted. The planar embedding can also be changed locally by flipping components that are connected to the rest of the graph by at most two vertices. Given vertices u, v, linkable(u, v) decides whether u and v are linkable in the current embedding, and if so, returns a list of suggestions for the placement of (u, v) in the embedding. For non-linkable vertices u, v, we define a new query, one-flip-linkable(u, v) providing a suggestion for a flip that will make them linkable if one exists. We support all updates and queries in $\mathcal{O}(\log^2 n)$ time. Our time bounds match those of Italiano et al. for a static (flipless) embedding of a dynamic graph. Our new algorithm is simpler, exploiting that the complement of a spanning tree of a connected plane graph is a spanning tree of the dual graph. The primal and dual trees are interpreted as having the same Euler tour, and a main idea of the new algorithm is an elegant interaction between top trees over the two trees via their common Euler tour.

1 Introduction

We present a data structure for supporting and maintaining a dynamic planar embedding of a dynamic graph. In this article, a *dynamic graph* is a graph where edges can be removed or inserted, and vertices can be cut or joined, but where an edge (u, v) can only be added if it does not violate planarity. More precisely, the edges around each vertex are ordered cyclically by the embedding, similar to the edge-list representation of the graph. A *corner* (of a face) is the gap between consecutive edges incident to some vertex. Given two corners c_u and c_v of the same face f, incident to the vertices u and v respectively, the operation insert(c_u, c_v) inserts an edge between u and v in the dynamic graph, and embeds it across f via the specified corners. We provide an operation linkable(u, v) that returns such a pair of corners c_u and c_v if they exist. If there are more options, we can list them in constant time per option after the first. A vertex may be cut through two corners, and linkable vertices may be joined by corners incident to the same face, if they are connected, or incident to any face otherwise. That is, joining vertices corresponds to linking them across a face with some edge e, and then contracting e.

It may often be relevant to change the embedding, e.g. in order to be able to insert an edge. In a *dynamic embedding*, the user is allowed to change the embedding by what we call flips, that is, to turn part of the graph upside down in the embedding. Of course, the relevance of this depends on what we want to describe with a dynamic plane graph. If the application is to describe roads on the ground, flipping orientation would not make much sense. But if we have the application of graph drawing or chip design in mind, flips are indeed relevant. In the case of chip design, a layer of a chip is a planar embedded circuit, which can be thought of as a planar embedded graph. An operation similar to flip is also supported by most drawing software.

Given two vertices u, v, we may ask whether they can be linked after modifying the embedding with only one flip. We introduce a new operation, the one-flip-linkable(u, v) query, which answers that question, and returns the vertices and corners describing the flip if it exists.

Our data structure is an extension to a well-known duality-based dynamic representation of a planar embedded graph known as a tree/cotree decomposition [4]. We maintain top trees [1] both for the primal and dual spanning trees. We use the fact that they share a common (extended) Euler tour - in a new way - to coordinate the updates and enable queries that either tree could not answer by itself. All updates and queries in the combined structure are supported in time $\mathcal{O}(\log^2 n)$, plus, in case of linkable(u, v), the

length of the returned list.

1.1 Dynamic Decision Support Systems

An interesting and related problem is that of dynamic planarity testing of graphs. That is, we have a planar graph, we insert some edge, and ask if the graph is still planar, that is, if there exist an embedding of it in the plane?

The problem of dynamic planarity testing appears technically harder than our problem, and in its basic form it is only relevant when the user is completely indifferent to the actual embedding of the graph. What we provide here falls more in the category of a decision support system for the common situation where the desired solution is not completely captured by the simple mathematical objective, in this case planarity. We are supporting the user in finding a good embedding, e.g., informing the user what the options are for inserting an edge (the linkable query), but leaving the actual choice to the user. We also support the user in changing their mind about the embedding, e.g. by flipping components, so as to make edge insertions possible. Using the one-flip-linkable query we can even suggest a flip that would make a desired edge insertion possible if one exists.

1.2 Previous work

Dynamic graphs have been studied for several decades. Usually, a *fully dynamic* graph is a graph that may be updated by the deletion or insertion of an edge, while *decremental* or *incremental* refer to graphs where edges may only be deleted or inserted, respectively. A dynamic graph can also be one where vertices may be deleted along with all their incident edges, or some combination of edge- and vertex updates [15].

Hopcroft and Tarjan [9] were the first to solve planarity testing of static graphs in linear time. Incremental planarity testing was solved by La Poutré [13], who improved on work by Di Battista, Tamassia, and Westbrook [2, 3, 19], to obtain a total running time of $\mathcal{O}(\alpha(q, n))$ where q is the number of operations, and where α is the inverse-Ackermann function. Galil, Italiano, and Sarnak [8] made a data structure for fully dynamic planarity testing with $\mathcal{O}(n^{2/3})$ worst case time per update, which was improved to $\mathcal{O}(n^{1/2})$ by Eppstein et al. [5]. For maintaining embeddings of planar graphs, Italiano, La Poutré, and Rauch [10] present a data structure for maintaining a planar embedding of a dynamic graph, with time $\mathcal{O}(\log^2 n)$ for update and for linkable-query, where insert only works when compatible with the embedding. The dynamic tree/cotree decomposition was first introduced

by Eppstein et al. [6] who used it to maintain the Minimum Spanning Tree (MST) of a planar embedded dynamic graph subject to a sequence of change-weight$(e, \Delta x)$ operations in $\mathcal{O}(\log n)$ time per update. Eppstein [4] presents a data structure for maintaining the MST of a dynamic graph, which handles updates in $\mathcal{O}(\log n)$ time if the graph remains plane. More precisely the user specifies a combinatorial embedding in terms of a cyclic ordering of the edges around each vertex. Planarity of the user specified embedding is checked using Euler's formula. Like our algorithm, Eppstein's supports flips. The fundamental difference is that Eppstein does not offer any support for keeping the embedding planar, e.g., to answer linkable(u, v), the user would in principle have to try all n^2 possible corner pairs c_u and c_v incident to u and v, and ask if insert(c_u, c_v) violates planarity.

As far as we know, the one-flip-linkable query has not been studied before. Technically it is the most challenging operation supported in this paper.

The highest lower bound for the problem of planarity testing is Pătraşcu's $\Omega(\log n)$ lower bound for fully dynamic planarity testing [14]. From the reduction it is clear that this lower bound holds as well for maintaining an embedding of a planar graph as we do in this article.

1.3 Results

We present a data structure for maintaining an embedding of a planar graph under the following updates:

delete(e) Deletes the specified edge, e.

insert(c_u,c_v) Inserts an edge through the specified corners, if compatible with the current embedding.

linkable(u,v) Returns a data structure representing the (possibly empty) list of faces that u and v have in common, and for each common face, f, the (nonempty) lists of corners incident to u and v, respectively, on f.

articulation-flip(c_1,c_2,c_3) If the corners c_1, c_2, c_3 are all incident to the same vertex, and c_1, c_2 are incident to the same face, this update flips the subgraph between c_1 and c_2 to the space indicated by c_3.

separation-flip(c_1, c_2, c_3, c_4) If the corners c_1, c_2 are incident to the vertex v, and c_3, c_4 are incident to the vertex u, and c_2, c_3 are incident to the face f, and c_1, c_4 are incident to the face g. Flips the subgraph delimited by these corners, or, equivalently, by u, g, v, f.

one-flip-linkable(u,v) Returns a suggestion for an articulation-flip or separation-flip making u and v linkable, if they are not already.

1.4 Paper outline.

In Section 2, we give first a short introduction to the data structures needed in our implementation, then we describe the implementations of the functions linkable (Section 2.3), insert and delete (Section 2.4), and articulation-flip and separation-flip (Section 2.5). Finally, Section 3 is dedicated to our implementation of one-flip-linkable.

2 Maintaining a dynamic embedding

In this section we present a data structure to maintain a dynamic embedding of a planar graph. In the following, unless otherwise stated, we will assume $G = (V, E)$ is a planar graph with a given combinatorial embedding and that $G^* = (F, E^*)$ is its dual.

Our primary goal is to be able to answer linkable(u, v), where u and v are vertices: Determine if an edge between u and v can be added to G without violating planarity and without changing the embedding. If it can be inserted, return the list of pairs of *corner*s (see Definition 1 below). Each corner-pair, (c_u, c_v), describes a unique place where such an edge may be inserted. If no such pair exists, return the empty list.

The data structure must allow efficient updates such as insert, remove, cut, join, and flip. We defer the exact definitions of these operations to Section 2.4.

As in most other dynamic graph algorithms we will be using a spanning tree as the main data structure, and note:

Observation 1. *[Tutte [18, p. 289]] If $E_T \subseteq E$ induces a spanning tree T in G, then $(E \setminus E_T)^*$ induces a spanning tree $\overline{T}^*$ in G^* called the* co-tree *of T.*

Lemma 1. *If u and v are vertices, T is a spanning tree, and e is any edge on the T-path between u and v, then any face containing both u and v lies on the fundamental cycle in G^* induced by e^* in the co-tree $\overline{T}^*$.*

Proof. The fundamental cycle $C(e)$ induces a cut of the graph which separates u and v.

Let f be a face which is incident to both u and v. Then there will be two paths between u and v on the face f, choose one of them, p. Since $C(e)$

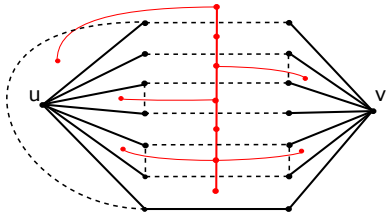

Figure 1: The co-tree path may switch arbitrarily between faces that are adjacent to any combination of neither, one, or both of u, v.

induces a cut of the graph which separates u and v, some edge $e_C \in p$ must belong to the cut. But then, f is a dual endpoint of e_C^*, and thus, f belongs to $C(e)$. □ □

Thus, the main idea is to search a path in the co-tree. This is complicated by the fact (see Figure 1) that the set of faces that are adjacent to u and/or to v need not be contiguous in $\overline{T}^*$, so it is possible for the cycle to change arbitrarily between faces that are adjacent to any combination of neither, one, or both of u, v.

Our linkable query will consist of two phases. A marking phase, in which we "activate" (mark) all corners incident to each of the two vertices we want to link (see Section 2.2), and a searching phase, in which we search for faces incident to "active" (marked) corners at both vertices (see Section 2.3). But first, we define corners.

2.1 Corners and the extended Euler tour

The concept of a corner in an embedded graph turns out to be very important for our data structure. Intuitively it is simply a place in the edge list of a vertex where you might insert a new edge, but as we shall see it has many other uses.

Definition 1. If G is a non-trivial connected, combinatorially embedded graph, a *corner* in G is a 4-tuple (f, v, e_1, e_2) where f is a face, v is a vertex, and e_1, e_2 are edges (not necessarily distinct) that are consecutive in the edge lists of both v and f.

If $G = (\{v\}, \emptyset)$ and $G^* = (\{f\}, \emptyset)$, there is only one corner, namely the tuple $(f, v, \mathbf{nil}, \mathbf{nil})$.

Note that faces and vertices appear symmetrically in the definition. Thus, there is a one-to-one correspondence between the corners of G and the corners of G^*. This is important because it lets us work with corners in G and G^* interchangeably. Another symmetric structure is the *Extended Euler Tour* (defined in Definition 2 below). Recall that an Euler Tour of a tree

is constructed by doubling all edges and finding an Eulerian Circuit of the resulting multigraph. This is extended in the following the following way:

Definition 2. Given a spanning tree T of a plane graph G, an *oriented* Euler tour is one where consecutive tree edges in the tour are consecutive in the ordering around their shared vertex.

Given the oriented Euler tour of the spanning tree T of a non-trivial connected graph G, we may define the *extended Euler tour* $\mathrm{EET}(T)$ by expanding between each consecutive pair of edges (e_-, e_+) with the list of corners and non-tree edges that come between e_- and e_+ in the orientation around their shared vertex. If G is a single-vertex graph, its extended Euler tour consist solely of the unique corner of G.

Thus, $\mathrm{EET}(T)$ is a cyclic arrangement that contains each edge in G exactly twice and each corner in G exactly once (see Figure 3). Even more interesting is:

Lemma 2. *Given a spanning tree T of G with extended Euler tour $\mathrm{EET}(T)$, the corresponding tour $\mathrm{EET}(\overline{T}^*)$ in G^* defines exactly the opposite cyclical arrangement of the corresponding edges and corners in G^*.*

In order to prove the lemma, we use the following definition of an *Euler cut*:

Definition 3. Given any closed simple curve C on the sphere S^2, then $S^2 \setminus C$ has two components, A and B, each homeomorphic to the plane[1]. Given an edge e of the embedded connected graph G with dual graph G^*, and a closed simple curve C, we say that C intersects e if C intersects either e of G, or e^* of G^*. Given a tree/cotree decomposition for G, we call C an *Euler cut* if the tree is entirely contained in A and the cotree is entirely contained in B, and the curve only intersects each edge, e or e^*, twice.

Proof of Lemma 2. Given any Euler cut C, note that the edges visited along the Euler tour of G come in exactly the same order as they are intersected by C, if we choose the same ordering of the plane (clockwise or counterclockwise). The placement of the corners is uniquely defined by the ordering of the edges.

Such an Euler cut always exists, because both tree and cotree are connected, and because they share no common point. One can construct it by contracting one of the two trees to a point, and then considering an ε-circle B_ε around that point, where ε is small enough such that there is a minimal

[1] Jordan curve Theorem.

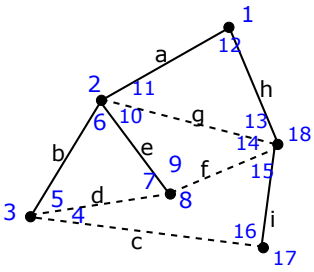

Figure 3: This graph has extended Euler tour `1a2b3c4d5b6e7d8f9e...18h`, or, to write only the edges: **ab**cd**be**df**eg**a**h**gf**i**c**ih**. Edges not in the spanning tree are drawn with dotted lines.

number of edge-intersections with that circle. The pre-image of B_ε under the contraction will be an Euler cut. (See Figure 2.)

Now, the Euler tour of G is uniquely given by C. But we also know that G is the dual of G^*, so C is also an Euler cut for G^*, and thus, the Euler tour of G^* is uniquely given by C.

Lastly, notice that in order for C to determine the same ordering of the edges, we need to orient the two components oppositely, that is, one clockwise and one counter-clockwise. □ □

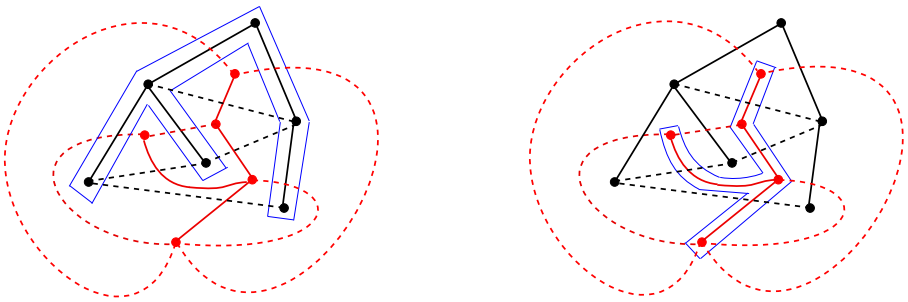

Figure 2: Two Euler cuts (blue) separating the primal tree from the dual tree.

Thus, segments of the extended Euler tour translate directly between T and $\overline{T}^*$. By *segment*, we mean any contiguous sub-list of the cycle.

The high level algorithm for linkable(u, v), that is, to find a face in the embedding incident to both u and v, (to be explained in detail later,) is now to build a structure consisting of an arbitrary spanning tree T and its co-tree $\overline{T}^*$ such that we can

1. Find an edge e on the T-path between u and v. This is easy.
2. Mark all corners incident to u and v in T. This is complicated by the existence of vertices of high degree, so a lazy marking scheme is

needed. However, it is easier than marking them in $\overline{T}^*$ directly, since each vertex has a unique place in T and no place in $\overline{T}^*$.

3. Transfer those marks from T to $\overline{T}^*$ using Lemma 2. We can do this as long as the lazy marking scheme works in terms of segments of the extended Euler tour.
4. Search the fundamental cycle induced by e^* in $\overline{T}^*$ for faces that are incident to a marked corner on both sides of the path.

2.2 Marking scheme

We need to be able to mark all corners incident to the two query vertices u and v, and we need to do it in a way that operates on segments of the extended Euler tour. To this end we will maintain a modified version of a top tree over T (see Alstrup et al. [1] for an introduction to top trees), called the extended embedded top tree.

2.2.1 Embedded top trees.

Given a tree T, a top tree for T is a binary tree. At the lowest level, its leaves are the edges of T. Its internal nodes, called *clusters*, are sub-trees of T with at most two boundary vertices. The set of boundary vertices of a cluster C is denoted ∂C. At the highest level its root is a single cluster containing all of T. A non-boundary vertex of the subset $S \subset V$ may not be adjacent to a vertex of $V \setminus S$. A cluster with two boundary vertices is called a *path cluster*, and other clusters are called *leaf clusters*. Any internal node is formed by the merged union of its (at most two) children. All operations on the top tree are implemented by first breaking down part of the old tree from the top with $\mathcal{O}(\log n)$ calls to a *split* operation, end then building the new tree with $\mathcal{O}(\log n)$ calls to a *merge* operation. (This was proven by Alstrup et al. in [1].)

The operation *expose* on the top tree takes one or two vertices and makes them boundary of the root cluster. We may even expose any constant number of vertices, and then, instead of having one top tree with the cluster containing all of T in the root, we obtain a constant number of top trees with those vertices exposed.

Definition 4. An *embedded top tree* is a top tree over an embedded tree T with the additional property that whenever it is merging $\{C_i\}$ to form C, for each pair C_i, C_{i+1} there exist edges $e_i \in C_i$ and $e_{i+1} \in C_{i+1}$ that are adjacent in the cyclic order around the common boundary vertex.

Definition 5. Given an embedded graph G and a spanning tree T of G, an *extended embedded top tree* of T is an embedded top tree over the embedded tree T' defined as follows:

- The vertices of T' are the vertices of G, plus one extra vertex for each corner of G and two extra vertices for each non-tree edge (one for each end vertex of the edge).
- The edges of T' are the edges of T, together with edges connecting each vertex v of G with the vertices representing the corners and non-tree edges incident to v.
- The cyclic order around each vertex is the order inherited from the embedding of G.

Observation 2. *In any extended embedded top tree over a tree T, path clusters consist of the edges and corners from two segments of* $\mathrm{EET}(T)$*, and leaf clusters consist of the edges and corners from one segment of* $\mathrm{EET}(T)$*.*

Lemma 3. *We can maintain an extended embedded top tree over T using $\mathcal{O}(\log n)$ calls to merge and split per update.*

Proof. First note that we can easily maintain T'. Now construct the ternary tree $\widehat{T}'$ by substituting each vertex by its chain (see Italiano et al. [10]). Consider the top tree for T' made via reduction to the topology tree for $\widehat{T}'$ (see Alstrup et al. [1]). Since the chain is defined to have the vertices in the same order as the embedding, this top tree is naturally an embedded top tree. □ □

For the rest of this paper, we will assume all top trees are extended embedded top trees, without any further mention.

Observation 3. *Since corners in T are represented by vertices in T', we automatically get the ability to* expose *corners, giving complete control over which* $\mathrm{EET}(T)$ *segment is available for information or modification.*

2.2.2 Slim-path top trees and four-way merges.

We can tweak the top tree further such that the following property, called *the slim path invariant*, is maintained:

- for any path-cluster, all edges in the cluster that are incident to a boundary vertex belong to the cluster path. In other words, for each boundary vertex b, there is at most one edge in the cluster that is incident to b.

This can be done by allowing the following merge: Given path-clusters C, C' with boundary vertices $\partial C = \{u, v\}$ and $\partial C' = \{v, w\}$, and given two leaf-clusters with boundary vertex v, allow the merge of those four, to obtain one path cluster with boundary vertices $\{u, w\}$. For lack of a better word, we call this a *four-way merge*.

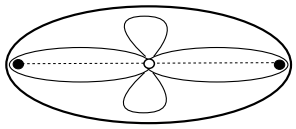

Figure 4: A four-way merge of 2 paths and 2 leaves.

We allow the following merges:

- Two leaves merge to one,
- A path cluster with boundary vertices u, v, and a leaf cluster with boundary vertex v, merge to form a leaf with boundary vertex u,
- Two path clusters with boundary vertices u, v, and v, w, respectively, and up to two leaves with boundary vertex v make a four-way merge to form a path cluster with boundary vertices u, v (see Figure 4).

Since the slim-path invariant has no restriction on leaves, clearly the first two do not break the invariant. The four-way merge does not break the invariant: If, before the merge, there was only one edge incident to u and only one incident to w, this is still maintained after the merge.

Definition 6. A *slim-path top tree* is a top tree which maintains the slim-path invariant by allowing four-way merges.

Lemma 4. *We can maintain a slim-path top tree over a dynamic tree with height* $\mathcal{O}(\log n)$*, using* $\mathcal{O}(\log n)$ *calls to* merge *and* split *per update.*

Proof. Proof by reduction to top trees (see figure 5). For each path cluster in the ordinary top tree, keep track of a slim path cluster, and up to four internal leaves; two for each boundary vertex, one to each side of the cluster path. In the slim-path top tree, this would correspond to four different clusters, but this doesn't harm the asymptotic height of the tree. When merging two path clusters in the ordinary top tree, with paths $u \cdots v$ and $v \cdots w$, this simply correspond to two levels of merges in the slim-path top tree: First, the incident leaf clusters are merged, pairwise, and then, we perform a four-way merge. Other merges in the ordinary top tree correspond exactly to

merges in the slim-path top tree. Thus, each merge in the ordinary top tree corresponds to at most 2 levels of merges in the slim-path top tree, and the maintained height is $\mathcal{O}(\log n)$. □ □

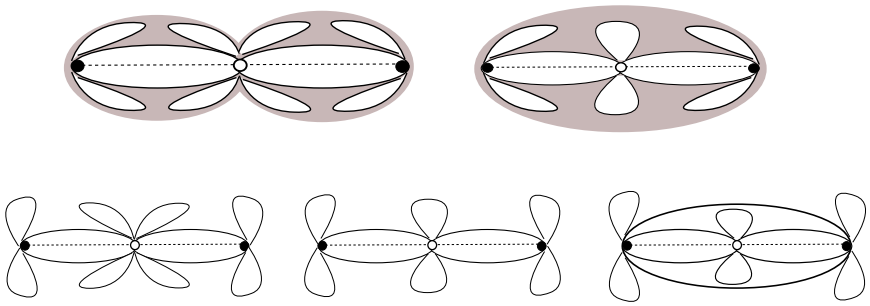

Figure 5: One merge in ordinary top trees corresponds to two merges in slim-path top trees.

Note that four-way merges and the slim-path invariant are not necessary. One could simply maintain, within each path cluster, information about the internal leaf clusters incident to each boundary vertex; one to each side of the cluster path. We present four-way merges because they provide some intuition about this particular form of usage of top trees.

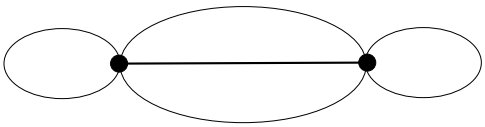

Figure 6: In a slim-path top tree, expose(u, v) exposes a path and two leaves.

We also note that the expose(u, v)-operation is different in a slim-path top tree, whenever u or v is not a leaf in the tree (see Figure 6.) Instead of returning a top tree which has at its top a path cluster with path $u \cdots v$, the operation returns up to three top trees: one with at its root a path cluster with path $u \cdots u$, and, up to two top trees with leaf clusters at their roots, one for each exposed non-leaf vertex.

Figure 7: The vertex w is a boundary vertex of the green clusters, but not of their blue parent clusters.

Lemma 5. *Whenever a merge of two clusters in the slim-path top tree causes a vertex w to stop being a boundary vertex (see Figure 7), all corners incident to w are contained in one or two* $\mathrm{EET}(T)$ *segments. These segments will be sub-segments of the (one or two) segments corresponding to the parent cluster C (blue in Figure 7), and will not contain any corners incident to the (one or two) boundary vertices of C.*

Proof. There are two statements to prove. Firstly, all corners incident to w are contained in the one or two $\mathrm{EET}(T)$ segments described in Observation 2. Secondly, it follows from the slim-path invariant that the delimiting corners for any boundary vertex $v \in \partial C$ are exactly those incident to the unique tree-edge in C incident to v. □ □

The following term is defined for both top trees and slim-path top trees and is very convenient:

Definition 7. Given a vertex v, let C_v denote the deepest cluster in the top tree where v is not a boundary vertex. The *root path* of v in the top tree is the path from C_v to the root in the top tree. For a corner c incident to a vertex v, the *root path* of c is the root path of v.

2.2.3 Deactivation counts

Now suppose we associate a (lazy) *deactivation count* with each corner. The deactivation count is set to 0 before we start building the top tree. Define the merge operation on the top tree such that whenever a merge discards a boundary vertex we *deactivate* all corners on the at most two segments of $\mathrm{EET}(T)$ mentioned in Lemma 5 by increasing that count (and define the split operation on the top tree to reactivate them as necessary). When the top tree is complete, the corners that are still *active* (have deactivation count 0) are exactly those incident to the boundary vertices of the root of the top tree. These boundary vertices are controlled by the expose operation on the top tree and changing the boundary vertices require only $\mathcal{O}(\log n)$ merges and splits, so we have now argued the following:

Lemma 6. *We can mark/unmark all corners incident to vertices u and v by increasing and decreasing the deactivation counts on $\mathcal{O}(\log n)$ segments of the extended Euler tour.*

What we really want, is to be able to search for the marked corners in $\overline{T}^*$, so instead of storing the counts (even lazily) in the top tree over T, we will store them in a top tree over $\overline{T}^*$. Again, each cluster in this top tree covers one or two segments of the extended Euler tour. For each segment S we keep track of a number $\Delta(S)$ such that for every corner, c, the deactivation count for c equals the sum of $\Delta(S)$ over all segments S containing c in the root path of c .

$$\text{deactivation count}(c) = \sum_{\substack{\text{cluster } C \in \text{root path}(c), \\ \text{segment } S \in C \,|\, c \in S}} \Delta(S)$$

In addition, for each boundary vertex (face) $f \in \partial C$, we keep track the value $c_{\min}(C, f)$, defined as the minimum deactivation count over all corners in the cluster C that are incident to f.

To update the deactivation counts of an arbitrary segment S, all we need to do is modify the $\mathcal{O}(\log n)$ clusters that are affected, which can be done in $\mathcal{O}(\log n)$ time, leading to

Lemma 7. *We can maintain a top tree over $\overline{T}^*$ that has $c_{\min}(C, f)$ for each boundary vertex $f \in \partial C$ where C is a cluster in $\overline{T}^*$, in $\mathcal{O}(\log^2 n)$ time per change in the set of (at most two) marked vertices.*

Proof. To mark a given set of vertices, expose them in the top tree for T. For each merge of clusters $\{C_i\}$ into a new cluster C in the top tree over T, where $\partial C \neq \bigcup \partial C_i$, consider the at most two segments of $\text{EET}(T)$ associated with C. For each such segment, S, expose the delimiting corners of S in the top tree over $\overline{T}^*$ (see Observation 3), and increase $\Delta(R)$ for the relevant segment R of the root cluster (in the dual top tree).

For each merge in the dual top tree, where $\{C_i\}$ merge to form C, and for each boundary vertex $f \in \partial C$, set $c_{\min}(C, f) = \min\{c_{\min}(C_i, f) \mid f \in \partial C_i\}$.

For each split of a cluster C into clusters $\{C_i\}$, and each segment S in C corresponding to segments $\{S_i\}$ with S_i in C_i, propagate $\Delta(S)$ down by setting $\Delta(S_i) = \Delta(S_i) + \Delta(S)$.

Since expose in the primal top tree can be implemented with $\mathcal{O}(\log n)$ calls to merge and split, this will yield at most $\mathcal{O}(\log n)$ calls to expose in the dual top tree, and thus, in total, the time for marking a set of vertices is $\mathcal{O}(\log^2 n)$. □ □

Observation 4. *This is enough for, given a face f and a vertex u, checking whether f is incident to u in $\mathcal{O}(\log^2 n)$ time.*

Proof. Expose u in the primal top tree. This takes $\mathcal{O}(\log^2 n)$ time. Expose f in the dual top tree. This takes $\mathcal{O}(\log n)$ time. Then, for the root cluster R which has boundary vertex f, $c_{\min}(R, f) = 0$ if an only if f and u are incident. □ □

2.3 Linkable query

Unfortunately, the $c_{\min}$ and Δ values discussed in Section 2.2 are not quite enough to let us find the corners we are looking for. We can use them to ask what marked corners a given face is incident to, but we do not have enough to find *pairs* of marked corners on opposite sides of the same face on the co-tree path.

As noted in Lemma 1, for two given vertices u and v, there exists a path in the dual tree containing all candidates for a common face. And this path is easily found! Since the dual of a primal tree edge induces a fundamental cycle that separates u and v, we may use the path between the dual endpoints f, g of any edge on the primal tree path between u and v. Furthermore, once we expose (f, g) in the dual tree, if $f \neq g$, the root will have two EET-segments: the minimum deactivation count of one EET-segment is 0 if and only if any non-endpoint faces are incident to v, the other is 0 if and only if any are incident to u. Checking the endpoint faces can be done (cf. Observation 4), but to find non-endpoint faces we need more structure.

To just output *one* common face, our solution is for each path cluster in the top tree over the co-tree to keep track of *a single* internal face $f_{\min}$ on the cluster path that is incident to minimally deactivated corners on either side of the cluster path if such a face exists. For that purpose, we define for any EET-segment S the value $c_{\min}(S)$ to represent the minimal deactivation count of a corner in S.

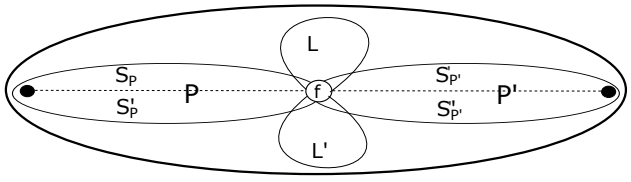

Figure 8: After a four-way cluster merge of L, L', P, P' to C, the face $f_{\min}(C)$, if it exists, is found among $f_{\min}(P)$, $f_{\min}(P')$, and f.

Lemma 8. *We can maintain a top tree over $\overline{T}^*$ that has $c_{\min}$ and Δ values for each EET-segment in each cluster and $f_{\min}$ values for each path cluster in $\mathcal{O}(\log^2 n)$ time per change to the set of (at most two) marked vertices.*

Proof. We maintain a slim-path top tree for the dual tree (see Lemma 4).

Merge and split of leaf-clusters are handled as in the proof of Lemma 7. When two leaf clusters L, L' with segments $S_L, S_{L'}$ merge, $c_{\min}$ for the resulting segment is set to be $\min\{c_{\min}(S_L), c_{\min}(S_{L'})\}$.

For a 4-way merge (see Figure 8) of two path clusters P, P' and two leaf clusters L, L' with common boundary vertex (face) f into a cluster C, we need to calculate $c_{\min}(S_C)$ and $c_{\min}(S'_C)$ for its two EET-segments, and the value of $f_{\min}(C)$. The EET-segment S_C is the concatenation of three EET-segments, S_P from P, $S_{p'}$ from P', and S_L from L. We may then set $c_{\min}(S_C) = \min\{c_{\min}(S_p), c_{\min}(S_{p'}), c_{\min}(S_L)\}$.

Similarly, $c_{\min}(S'_C) = \min\{c_{\min}(S'_p), c_{\min}(S'_{p'}), c_{\min}(S_{L'})\}$.

We may now assume we have calculated the values $c_{\min}(S_C)$ and $c_{\min}(S'_C)$. To determine $f_{\min}(C)$, one only has to check the three values: $f_{\min}(P), f_{\min}(P')$, and f. If $c_{\min}(S_C) = c_{\min}(S_P)$, and $c_{\min}(S'_C) = c_{\min}(S'_P)$, then we may set $f_{\min}(C) = f_{\min}(P)$. Otherwise, if $c_{\min}(S_C) = c_{\min}(S_{P'})$, and $c_{\min}(S'_C) = c_{\min}(S'_{P'})$, then we may set $f_{\min}(C) = f_{\min}(P')$. Otherwise, if $c_{\min}(S_C) = c_{\min}(L)$ and $c_{\min}(S'_C) = c_{\min}(L')$, we may set $f_{\min}(C) = f$. Otherwise, finally, we have to set $f_{\min}(C) = \textbf{nil}$. □ □

Lemma 9. *We can support each* linkable(u, v) *in $\mathcal{O}(\log^2 n)$ time per operation.*

Proof. If u and v are not in the same connected component we pick any corners c_u and c_v adjacent to u and v and return them. Otherwise, we use expose(u, v) on the top tree over T to activate all corners adjacent to u and v and to find an edge e on the T-path from u to v (e.g. the first edge on the path). Let g, h be the endpoints of e^*, and call expose(g, h) on the top tree over $\overline{T}^*$. By Lemma 1, a common face lies on the cotree path from g to h. Let f be the $f_{\min}$ value of the resulting root. We can now test each of f, g, h using the $c_{\min}$ values to find the desired corners if they exist. □ □

Lemma 10. *If there are more valid answers to* linkable(u, v) *we can find k of them in $\mathcal{O}(\log^2 n + k)$ time.*

Proof. For each leaf cluster and for each side of each path cluster we can maintain the list of minimally deactivated corners adjacent to each boundary vertex. Then, instead of maintaining a single face $f_{\min}$ for each path cluster, we can maintain a linked list of all relevant faces in the same time. And for

each side of each face in the list we can point to a list of minimally deactivated corners that are adjacent to that side. For leaf-clusters, we point to a linked list of minimally deactivated corners incident to the boundary vertex. Upon the merge of clusters, face-lists and corner-lists may be linked together, and the point of concatenation is stored in the resulting merged cluster in case of a future split. Note that each face occurs in exactly one face-list.

As before, to perform linkable(u, v), expose u, v in the primal tree. Let e_0 be an edge on the tree-path between u and v, and expose the endpoints of e_0^* in the dual top tree. Now, the maintained face-list in the root of the dual top tree contains all faces incident to u, v, except maybe the endpoints of e_0^*, which can be handled separately, as before. The total time is therefore $\mathcal{O}(\log^2 n)$ for the necessary expose operations, and then $\mathcal{O}(1)$ for each reply. □ □

Observation 5. *If we separately maintain a version of this data structure for the dual graph, then for faces f, g,* linkable(f, g) *in that structure lets us find vertices that are incident to both f and g.*

2.4 Updates

In addition to the query, our data structure supports the following set of update operations:

- insert(c_u, c_v) where c_u and c_v are corners that are either in different connected components, or incident to the same face. Adds a new edge to the graph, inserting it between the edges of c_u at one end and between the edges of c_v at the other. Returns the new edge.
- remove(e). Removes the edge e from the graph. Returns the two corners that could be used to insert the edge again.
- join(c_u, c_v) where c_u and c_v are corners that are either in separate components of the graph or in the same face. Combines the vertices u and v into a single new vertex w and returns the two new corners c_w and c'_w that may be used to split it again using cut(c_w, c'_w).
- cut(c_w, c'_w) where c_w and c'_w are corners sharing a vertex w. Splits the vertex into two new vertices u and v and returns corners c_u and c_w that might be used to join them again using join(c_u, c_v).

When calling remove(e) on a non-bridge tree-edge e, we need to search for a replacement edge. Luckily, e^* induces a cycle in the dual tree, and any other edge on that cycle is a candidate for a replacement edge. If we like, we can augment the dual top tree so we can find the minimal-weight replacement edge, simply let each path cluster remember the cheapest edge

on the tree-path, and expose the endpoints of e^*. If we want to keep T as a minimum spanning tree, we also need to check at each insert and join that we remove the maximum-weight edge on the induced cycle from the spanning tree.

In general, when we need to update both the top trees over T and $\overline{T}^*$ we must be careful that we first do the splits needed in the top tree over T to make each unchanged sub-tree a (partial) top tree by itself, then update the top tree over $\overline{T}^*$ and finally do the remaining splits and merges to rebuild the top tree over T. This is necessary because the merge and split we use for T depend on T and $\overline{T}^*$ having related extended Euler tours.

Any change to the graph, especially to the spanning tree, implies a change to the extended Euler tour. Furthermore, any deletion or insertion of an edge implies a merge or split in the dual tree. E.g. if an edge is inserted across a face, that face is split in two. As a more complex example, if the non-bridge tree-edge $e = (u, v)$ is deleted, the replacement edge is removed from the dual tree, and the endpoints of e^* are merged.

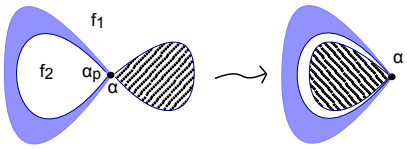

Figure 9: An articulation flip at the vertex α.

2.4.1 Cut and join of faces.

Given a face f and two corners incident to the face, say, c_1 and c_2, we may cut the face in those corners, producing new faces f and f'. We go about this by exposing c_1, c_2 in the dual top tree. We then have two chains for f, which we rename as f and f'.

Similarly, faces may be joined. Given a face f and a face g, a corner c_f incident to f and a corner c_g incident to g, we may join the two faces to a new face h, with two new corners c_1 and c_2, such that all edges incident to h between c_1 and c_2 are exactly those from f, in the same order, and those between c_2 and c_1 are those from g. We do this by exposing f and c_f in one top tree, and g and c_g in the other top tree, and then by the link and join operation.

2.4.2 Inserting an edge

We now show how to perform insert(c_u, c_v), when compatible with the embedding.

To insert an edge between two components, we must add it to the primal

tree. Given a vertex v with an incident corner c_v to the face g, and a vertex u with an incident corner c_u to the face f, we may insert an edge a between u and v.

In each primal top tree, we expose v, c_v and u, c_u, respectively. In the dual top tree, we expose g, c_v and f, c_u, respectively. We then link u and v, and update the EET correspondingly: New corners $c_1, \ldots, c_4$ are formed, such that the new edge $a = (u, v)$ appears as the successor of c_4 and the predecessor of c_1, and the successor of c_2 and the predecessor of c_3 (see figure 10). That is, we may view the top tree in its current form as the top tree with an exposed path (u, v) which is trivial (consisting of one edge) and two leaves, one at each end.

In the dual tree, the faces f and g are simply joined at the incident corners c_u and c_v, respectively.

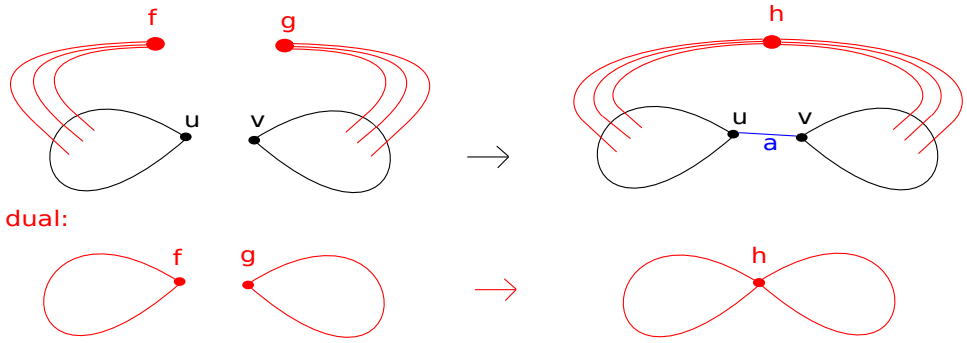

Figure 10: The component of u is joined with the component of v by the edge a, and new corners are formed.

To insert an edge inside a component, the two vertices must belong to the same face. Let u and v be vertices with corners c_u and c_v incident to a face f. We can then expose the corners c_u and c_v, which gives us a path cluster in the primal top tree (and some leaves, because we use slim-path top trees), and in the dual tree, exposing those corners gives us two leaves, both incident to f.

We now create new corners $c_1, \ldots, c_4$, and update the EET to list the new edge, $a = (u, v)$, as the successor of c_1 and predecessor of c_2, replacing c_v, and as the successor of c_3 and predecessor of c_4, replacing c_u (see figure 11). In the dual tree, the face f is split in the corners c_u and c_v, creating new faces g and h, and a new edge (g, h) is added to the dual tree.

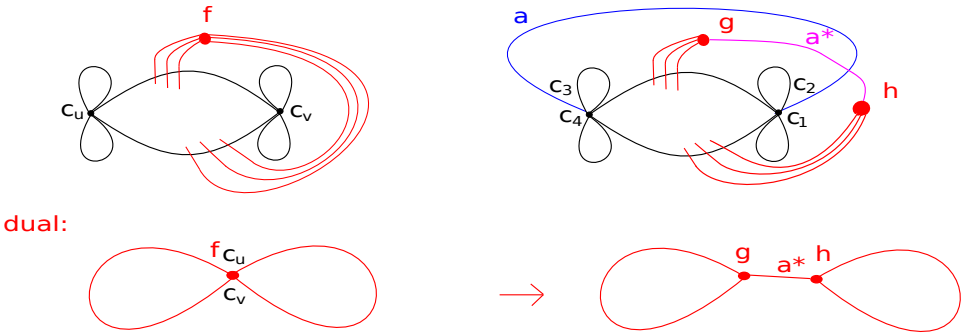

Figure 11: The corners c_u and c_v are exposed, and an edge $a = (u, v)$ (blue) is inserted in those corners. The dual of a, namely $a^* = (f, g)$ (pink), is inserted in the dual tree to connect the two new faces, that appear when the face f is cut in two by the insertion of the edge a.

2.4.3 Deleting an edge

We show now how to perform delete(e).

There are three cases, depending on whether the edge to be deleted is a bridge, a non-tree edge, or a replacement edge exists.

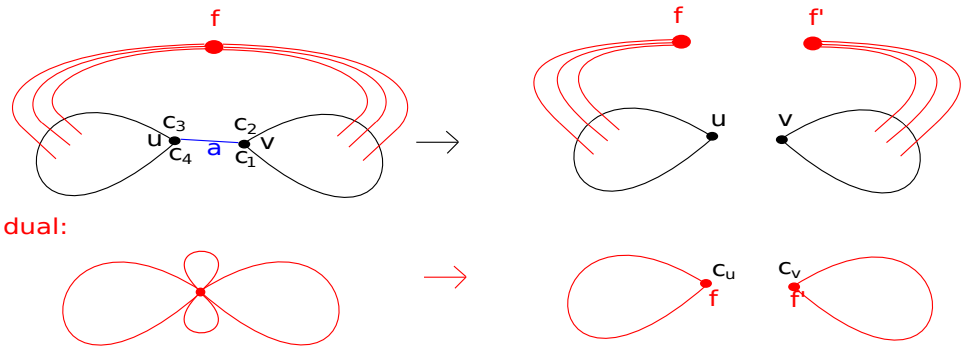

Figure 12: The bridge (u, v) is deleted.

An edge, e, is a bridge if its dual e^* is a self-loop of some face. To delete a bridge $e = (u, v)$ incident to the face f in corners $c_1, \ldots c_4$ (see figure 12), do the following. Expose u and v in the primal tree. Expose the corners $c_1, \ldots c_4$ in the dual tree, meaning we have at the top a four-leaf clover of clusters incident to f. Cut f into those four parts, and delete the leaves corresponding to $c_4 e c_1$ and $c_2 e c_3$. Let c_u denote the newly formed corner incident to u and c_v that incident to v. We have now split the face f in two,

and have two corresponding EETs and two dual top trees. In the primal top tree, delete the edge u; this leaves us with two top trees, one with a leaf with boundary vertex v exposed and one with boundary vertex u exposed – update the corresponding corners to be exactly the c_v and c_u created in the dual top tree, respectively.

If a non-bridge tree-edge is to be deleted, its dual induces a fundamental cycle in the dual tree, $C(e^*)$. Any edge on this cycle would reconnect the components and can be added to the primal tree as a replacement edge. For simplicity, we can just choose the edge after e^* on that cycle.

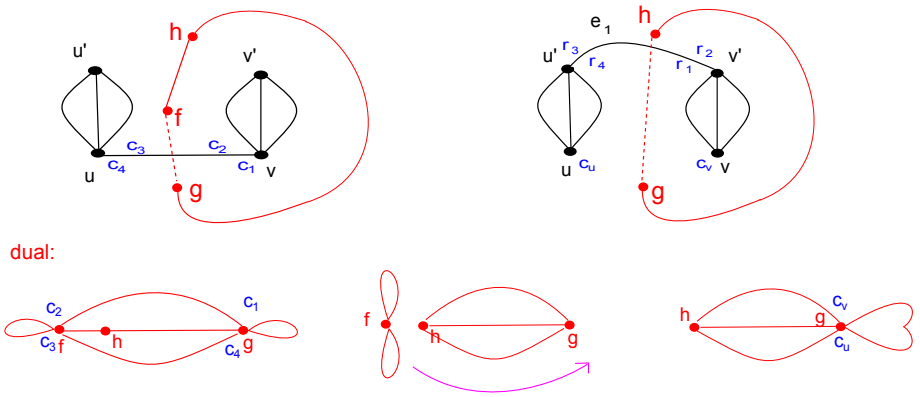

Figure 13: We delete the tree-edge $e = (u, v)$ from the graph. The non-tree edge (u', v') is added to the spanning tree, and the faces on either side of e are joined.

Let $e = (u, v)$ be the edge we want to delete. Let f and g be the faces incident to e. First, expose u and v in the primal tree. The top of the primal top tree is now of the form leaf-path-leaf, where the path-cluster consists of one edge, (u, v). Since $f \neq g$, exposing corners c_1, c_2 incident to e (and to v) in the dual tree returns a top tree with a path-cluster as root. Let $a_1 = (f, h) = e_1^*$ be the first edge on the cluster path from g to f. We want to replace e with $e_1 = (u', v')$. (See figure 13.)

Let c_h and c'_h be the corners of a_1 incident to h. Expose the vertices u, v, u', v' and the corners c_h and c'_h in the primal top tree.

Expose the faces f and g and the corners $c_1 \dots c_4$ in the dual top tree. In the dual tree, delete the edge $a_1 = (f, h)$ from the tree. Cut f in c_2, c_3 and g in c_1, c_4, delete the leaves corresponding to $c_2 e c_3$ and $c_4 e c_1$, and name the replacing corners c_f and c_g (updating the EET). Then, join the faces f and g in the corners c_f and c_g, and name the newly constructed corners c_u

and c_v. The dual top tree now has the path cluster with boundary vertices h and g as root.

In the primal tree, delete the edge (u, v). The newly formed corners are exactly the c_u and c_v formed in the dual top tree. Link the vertices u' and v', and update the EET correspondingly. That is, if r_1, r_2, r_3, r_4 are corners incident to $e_1 = (u', v')$, then exchange the segments $r_2 e_1 r_1$ and $r_4 e_1 r_3$ to $r_2 e_1 r_3$ and $r_4 e_1 r_1$.

Figure 14: A non-tree edge is exposed in the dual top tree, and removed.

Deleting a non-tree edge is the opposite of inserting an edge inside a face (see figure 11). If a non-tree edge $e = (u, v)$ is deleted, its incident faces f and g must be joined. In the primal tree, expose u, v and the corners $c_1, \ldots, c_4$ incident to e. In the dual tree, expose the edge $e^* = (f, g)$, which must belong to the cotree, and all four corners incident to e. Cut g in the corners c_1, c_4, cut f in the corners c_2, c_3. Delete the cluster containing e^*, and join f and g in the corners c_f and c_g which were created from the cut. The join produces new corners c_u and c_v. In the primal tree, the leaves with EET $c_2 e c_1$ and $c_4 e c_3$ are deleted, and the delimiting corners for the cluster path are updated to be the newly formed c_u and c_v. During this procedure, the EET is updated to contain c_v in place of $c_2 e c_1$ and c_u in place of $c_4 e c_3$.

2.5 Flip

Finally, for flip to work we have to use a version of top trees that is not tied to a specific clockwise orientation of the vertices. The version in [1] that is based on a reduction to Frederickson's topology trees [7] works fine for this purpose.

Definition 8 (Articulation flip)**.** Having cut and join functions, we may perform an *articulation-flip* (see Figure 15) — a flip in an articulation point: Given a vertex α incident to the face f_1 in two corners, c_1 and c_2, we may cut through c_1, c_2, obtaining two graphs G_1, G_2, having split α in vertices $\alpha_1 \in G_1, \alpha_2 \in G_2$, and having introduced new corners c, c' where we cut. Now, given a corner α_p incident to α_1 and incident to some face f_2, we may join α_1 with α_2 by the corners α_2, α_p, with or without having flipped the orientation of G_2.

Figure 15: A separation flip at a separation pair (blue). The flip makes vertex u linkable with vertex v.

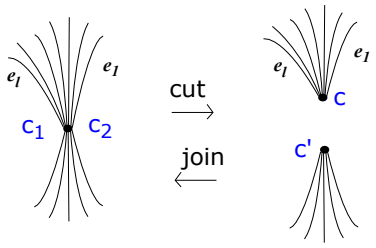

Figure 16: Left to right: The vertex is cut in two specified corners, and two new vertices are formed. Right to left: Two vertices are joined to one.

Definition 9 (Separation flip)**.** Similarly, given a separation pair α, β, incident to the faces f, g with corners $c_1, \ldots c_4$, we may cut through those corners, obtaining two graphs. We may then flip the orientation of one of them, and rejoin. We call this a *separation-flip*.

Internally, both flip operations are done by cutting out a subgraph, altering its orientation, and joining the subgraph back in. In the dual graph, a flip corresponds to two splits, two cuts, two links, and two merges.

First, the articulation flip will be described. To flip in an articulation point, we only need to be able to cut a vertex, and join two vertices to one. (See figure 16.)

2.5.1 Articulation flip

Given a vertex v and two corners c_1, c_2 incident to v, we may cut the vertex in two such that all edges between c_1 and c_2 belong to the new vertex v', and all edges between c_2 and c_1 belong to v. (See Figure 16.) The corners c_1 and c_2 are deleted, and new corners c and c' are created, and the extended Euler tours are updated correspondingly. (See figure 17.)

The EET cycle is split into two cycles; let e_1 and e_l denote the successor of c_2 and the predecessor of c_1, respectively, and let e'_1 and $e'_{l'}$ denote the successor of c_1 and the predecessor of c_2, respectively. The corner c will take place between e_l and e_1 in the EET containing v, and the corner c' will take place between $e'_{l'}$ and e'_1 in the EET containing v'. (See figure 17.)

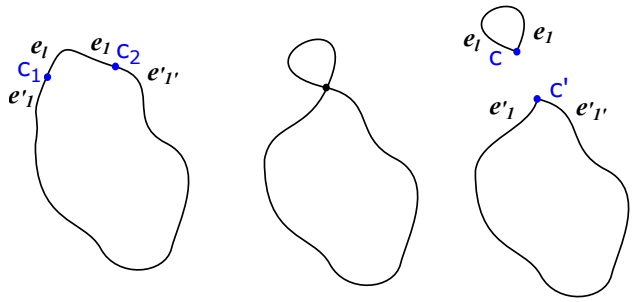

Figure 17: The extended Euler tour is cut in c_1 and c_2, and glued back together to form two tours. New corners, c and c' are formed.

We need to be able to cut both primal and dual vertices. By convention, the extended Euler tour is updated when primal vertices are cut, and never when dual vertices are cut.

Similarly, given two vertices, and given a corner incident to either vertex, we may join the two vertices, as an inverse operation to the vertex cut.

If c is a corner incident to the vertex v and c' is incident to v', we may perform the operation $\text{join}(c, c')$. Now, the extended Euler tour is updated correspondingly. We must create new corners, c_1 and c_2. In the Euler tour, c_1 must point to the predecessor of c as its predecessor, and the successor of c' as its successor, and similarly for c_2. And vice versa, the object of those edges must point to c_1 and c_2.

We now have the necessary tools for flipping in an articulation point, that is, if a vertex has two corners incident to the same face, we may flip the subgraph spanned by those corners into another face incident to the vertex. This is done by the two-step scheme of cut and join. Namely:

Theorem 1. *We can support* $\text{articulation-flip}(\mathrm{c}_1, \mathrm{c}_2, \mathrm{c}_3)$ *in time* $\mathcal{O}(\log^2 n)$.

Proof. Given a vertex v and two corners c_1 and c_2 incident both to v and some common face f_1, and given a third corner, c_3 incident to v and some face f_2, we first perform the vertex cut, $\text{cut}(v, c_1, c_2)$ in the primal top tree. Then, we perform the operation of $\text{cut}(f_1, c_1, c_2)$ in the dual top tree. Together, these two operations return two graphs, G and G', with respectively v, c and v', c' exposed in their primal top trees, and with f_1, c and f_1', c' exposed in their dual top trees. Then, we expose c_3 (incident to v and f_2) in the top

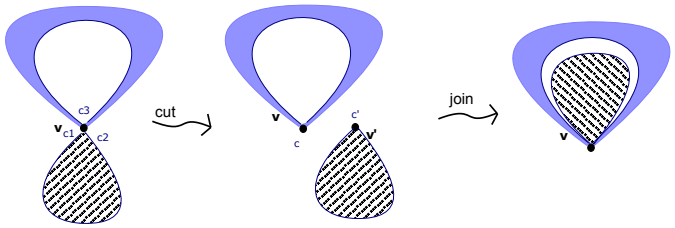

Figure 18: To flip in an articulation point: cut and join.

trees of G. Finally, we perform the operation join(c_3, c'), first in the dual, and then in the primal top tree, gluing G' back to G.

This operation takes $\mathcal{O}(\log^2 n)$ time, dominated by the expose-operation in the primal top tree. □ □

2.5.2 Altering the orientation

To perform a flip in a separation pair, we introduce the following three operations: seclude, alter the orientation, include. Seclude and include are similar to those of the previous section, while altering the orientation means that the ordering of edges around any vertex becomes the opposite. In the EET, the predecessor becomes a successor, and vice versa.

To implement this, we let the clusters of the primal and dual top tree contain one more piece of information, namely the orientation of the cluster, which is plus $+$, or minus $-$.

When a cluster is "negative", some of its information changes character:

- its left-hand child becomes a right-hand child and vice versa,
- if it is a path cluster, the two EET-segments switch places,
- and, finally, in the extended Euler tour, predecessor is interpreted as successor and vice versa.

When a cluster is split, its sign is multiplied with the signs of its children. That is, if a cluster with a minus is split, the minus is propagated down to the children. When clusters are merged, the new union-cluster is simply equipped with a plus, while its children clusters keep their sign.

In the 3-step program of "seclude, alter, include," the alter-step simply consists of changing the sign of the top cluster of the dual top tree for the secluded graph.

2.5.3 Separation flip

Given four corners $c_1, \ldots, c_4$ such that c_1, c_2 are incident to the vertex v, c_3, c_4 are incident to the vertex u, c_2, c_3 are incident to the face f and c_4, c_1 are incident to the face g. See figure 19.

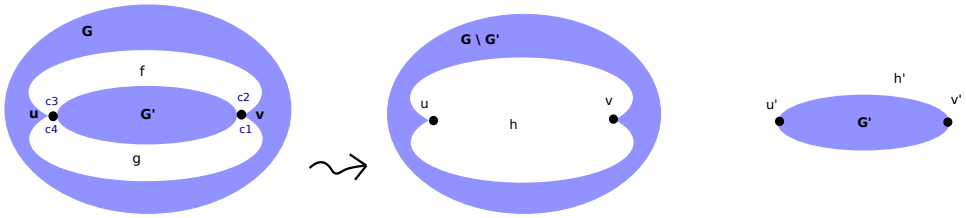

Figure 19: The subgraph G' is delimited by two vertices and two faces. The picture shows how G' is cut out of the graph G.

To seclude(c_1, c_2, c_3, c_4), cut the vertex v through the corners c_1 and c_2, cut u through c_3 and c_4. In the dual top tree, cut the face f in c_2 and c_3 (obtain f and f'), and cut the face g in c_4 and c_1 (which returns g and g'). Then, join f' with g', and f with g in the dual top tree. We now have two graph-stumps, denote them G and G', where G contains u and v which are both incident to one face h which was formed by joining f with g, and where G' contains u' and v', both incident to the new face h'.

This procedure consists of two vertex cuts in the primal tree, this takes $\mathcal{O}(\log^2 n)$ time, two vertex cuts in the dual tree, which takes the time $\mathcal{O}(\log n)$. Note that between the seclude and the subsequent include we may not have a spanning tree for one of the components. This is not a problem, since we are not doing any other operations in between.

Include is the inverse operation of seclude above.

Given two graphs, G and G', and given two vertices $u, v \in G$ and two vertices $u', v' \in G'$, and given a designated face $h \in G$ incident to both u and v, and similarly a face $h' \in G'$ incident to u' and v'. Let four corners be given, such that $c_{u,h}$ is incident to u and h, and so on, $c_{v,h}, c_{u',h'}, c_{v',h'}$.

To perform include$(c_{u,h}, c_{v,h}, c_{u',h'}, c_{v',h'})$: Cut the face h through the corners $c_{v,h}$ and $c_{v,h}$, and cut the face h' similarly. Let the resulting faces be denoted f, g, and f', g', respectively. Now, join the faces f with f' and g with g'. Finally, join the vertices u with u' and v with v'.

Theorem 2. *We can support* separation-flip(c_1, c_2, c_3, c_4) *in time* $\mathcal{O}(\log^2 n)$.

Proof. Given two vertices which are incident to two faces, and given four delimiting corners $c_1 \ldots c_4$ pairwise incident to both vertices and both faces

(as in Definition 9), we perform separation-flip(c_1, c_2, c_3, c_4) with a three-step procedure: Seclude, alter, include.

Let $g, f, c_1, \ldots, c_4$ be as in Definition 9 (see Figure 19). First, seclude the subgraph G' delimited by the corners $c_1, \ldots, c_4$, splitting the faces f and g to f, f' and g, g', respectively. Then, alter the orientation of G' by changing its sign. Finally, perform include, such that f joins with g', and g with f'.

This takes $\mathcal{O}(\log^2 n)$ time, dominated by the seclude and include operations. □ □

3 One-flip linkable query

Given vertices u, v, we have already presented a data structure to find a common face for u, v. Given they do not share a common face, we will determine if an articulation flip exists such that an edge between them can be inserted, and given no such articulation-flip exists, we will determine if a separation-flip that makes the edge insertion (u, v) possible exists.

Let f_1 and f_2 be faces in G, and let S be a subgraph of G. We say that S *separates* f_1 and f_2 if f_1 and f_2 are not connected in $G^* \setminus (E[S])^*$. Here, $E[X]$ denotes the set of edges of the subgraph X, $E[f]$ the edges incident to the face f, and $V[f]$ the incident vertices.

Observation 6. *Given a fundamental cycle C that is induced in $T \cup \{e\}$ by some edge e and given any two faces f_1, f_2 not separated by C, any face f such that $C \cup E[f]$ separates f_1 and f_2 lies on the path $f_1 \cdots f_2$ in $\overline{T}^*$.*

Let f_1 and f_2 be faces of G, and let $S = V[f_1] \cap V[f_2]$ be the set of vertices they have in common. Let C denote the set of corners between vertices in S and faces in $\{f_1, f_2\}$. The sub-graphs obtained by cutting G through all the corners of C are called *flip-components* of G w.r.t. f_1 and f_2. Flip-components which are only incident to one vertex of S can be flipped with an articulation-flip, and flip-components incident to two vertices can be flipped with a separation-flip. (See Figure 20.)

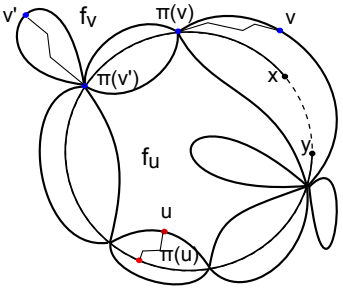

Figure 20: The faces f_u and f_v have five common vertices, and there are eight flip-components with respect to them.

Observation 7. *Note that the perimeter of a flip component always consists of the union of a path along the face of f_u with a path along the*

face of f_v. One of these paths is trivial (equal to a point) exactly when u, v are linkable via an articulation-flip.

Given vertices u, v in G, that are connected and not incident to a common face, we wish to find faces f_u and f_v such that u and v are in different flip-components w.r.t. f_u and f_v.

3.1 Finding one face

Let u and v be given vertices, and assume there exist faces f_u and f_v such that $u \in V[f_u] \setminus V[f_v]$, $v \in V[f_v] \setminus V[f_u]$, and u and v are in different flip-components w.r.t. f_u and f_v.

Let u_L, u_R be the left and right faces adjacent to the first edge on the path from u to v. Similarly let v_L, v_R be the left and right faces adjacent to the first edge on the path from v to u.

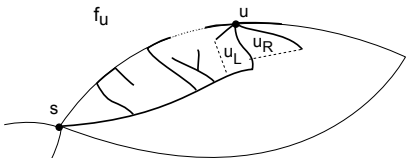

Figure 21: The co-tree path from u_L to u_R goes through f_u. The proof uses that the tree-path from u to v goes through some $s \in S$ on the boundary of u's flip-component.

Lemma 11. *Face f_u is on the $\overline{T}^*$-path $u_L \cdots u_R$ and face f_v is on the $\overline{T}^*$-path $v_L \cdots v_R$.*

Proof. For symmetry reasons, we need only be concerned with the case f_u. The perimeter of a flip-component consists of edges incident to f_u and edges incident to f_v (see Observation 7). Furthermore, in order for u, v to be linkable via a flip, u needs to lie on the perimeter of its flip-component. We also know that the tree-path from v to u must go through a point s in S which lies on the boundary of u's flip-component. Thus, there must exist a path p in G from $\pi \in S$ to u, consisting only of edges incident to f_u. Note that $u \notin S$ since u, v were not already linkable. If the first edge e_u on the tree path from u to v is not already incident to f_u, then the union of p and the tree must contain a fundamental cycle containing e_u, separating u_L from u_R, induced by an edge e_i incident to f_u. (See Figure 21.) But then, the co-tree path from u_L to u_R goes through e_i^*, which means it goes through f_u. □ □

Lemma 12. *If there exists a fundamental cycle C separating f_u from f_v such that $u \notin V[C]$ and $v \in V[C]$, then $f_u = \mathrm{m}(u_L, u_R, f)$ where $f \in \{v_L, v_R\}$ is the face that is on the same side of C as u. Here, $\mathrm{m}(a, b, c)$ denotes a's projection to the path $b \cdots c$.*

Proof. By Lemma 11 f_u is on the path $u_L \cdots u_R$. And since $C \cup E[f_u]$ separates f from u_L and u_R it is on the paths $u_L \cdots f$ and $u_R \cdots f$ by Observation 6. □ □

Lemma 13. *If there exists a fundamental cycle C separating f_u from f_v such that $u \notin V[C]$ and $v \notin V[C]$, then either $f_u = \mathrm{m}(u_L, u_R, v_L) = \mathrm{m}(u_L, u_R, v_R)$ or $f_v = \mathrm{m}(v_L, v_R, u_L) = \mathrm{m}(v_L, v_R, u_R)$.*

Proof. Let e be the edge in $C \setminus T$, and let e_u, e_v be the faces adjacent to e that are on the same side of C as f_u and f_v respectively. Then e is on all 4 paths in $\overline{T}^*$ with u_L or u_R at one end and v_L or v_R at the other. At least one of u, v is in a different flip-component from e, so we can assume without loss of generality that u is. By Lemma 11 f_u is on the path $u_L \cdots u_R$. And since $C \cup f_u$ separates u_L and u_R from e_u, f_u is on both the paths $u_L \cdots e_u$ and $u_R \cdots e_u$ by Observation 6. Thus $f_u = \mathrm{m}(u_L, u_R, e_u) = \mathrm{m}(u_L, u_R, v_L) = \mathrm{m}(u_L, u_R, v_R)$. □ □

Lemma 14. *If a fundamental cycle C separates f_u from f_v such that $u \in V[C]$ and $v \in V[C]$, then either $f_u = \mathrm{m}(u_L, v_L, v_R) = \mathrm{m}(u_L, u_R, v_R)$ or $f_u = \mathrm{m}(u_R, v_L, v_R) = \mathrm{m}(u_L, u_R, v_L)$ or $f_v = \mathrm{m}(v_L, u_L, u_R) = \mathrm{m}(v_L, v_R, u_R)$ or $f_v = \mathrm{m}(v_R, u_L, u_R) = \mathrm{m}(v_L, v_R, u_L)$.*

Proof. Let e be the edge in $C \setminus T$, and let e_u, e_v be the faces adjacent to e that are on the same side of C as f_u and f_v respectively. Then e is on all 4 paths in $\overline{T}^*$ with u_L or v_R at one end and v_L or u_R at the other. Assume that u_L and v_R are on the side of C containing f_u and u_R and v_L are on the side of C containing f_v. At least one of u, v is in a different flip-component from e, so assume that v is. By Lemma 11 f_u is on the path $u_L \cdots e_u \subset u_L \cdots u_R$. And since $C \cup f_u$ separates u_L and e_u from v_R it is on both the paths $u_L \cdots v_R$ and $e_u \cdots v_R$ by Observation 6. Thus $f_u = \mathrm{m}(u_L, e_u, v_R) = \mathrm{m}(u_L, v_L, v_R) = \mathrm{m}(u_L, u_R, v_R)$. The remaining cases are symmetric. □ □

Theorem 3. *If f_u, f_v exist, either $f_u \in \{\mathrm{m}(u_L, u_R, v_L), \mathrm{m}(u_L, u_R, v_R)\}$ or $f_v \in \{\mathrm{m}(u_L, v_L, v_R), \mathrm{m}(u_R, v_L, v_R)\}$.*

Proof. If they exist there is at least one fundamental cycle C separating them. This cycle must have the properties of at least one of Lemmas 12, 13, or 14. □ □

By computing the at most two different m values and checking which ones (if any) contain u or v we therefore get at most two candidates and are guaranteed that at least one of them is in $\{f_u, f_v\}$ if they exist.

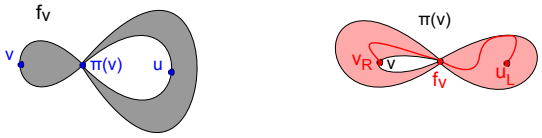

Figure 22: If f_v is an articulation point, so is $\pi(v)$. But then the co-tree path from u_L to v_L must go through f_v. Left: Primal graph. Right: Dual graph.

3.2 Finding the other face

Lemma 15. *Let u, v, and f_u be given. Then the first edge e_L on $f_u \cdots v_L$ or the first edge e_R on $f_u \cdots v_R$ induces a fundamental cycle $C(e_R)$ or $C(e_L)$ in T that separates f_u from f_v.*

Proof. By lemma 11, f_v is on $v_L \cdots v_R$ in $\overline{T}^*$, so the first edge on $f_u \cdots f_v$ is also the first edge on either $f_u \cdots v_L$ or $f_u \cdots v_R$. □ □

Thus given the correct f_u we can find at most two candidates for an edge e that induces a fundamental cycle $C(e)$ in T that separates f_u from f_v, and be guaranteed that one of them is correct.

Observation 8. *For each vertex, v, we may consider the projection $\pi(v)$ of v onto the cycle C. For each flip-component, X, we may consider the projection $\pi(X) = \{\pi(v) \mid v \in X\}$. If X is an articulation-flip component, the projection $\pi(X)$ is a single point in $S = V[f_u] \cap V[f_v]$. If X is a separation-flip component, its projection is a segment of the cycle, $\pi_1 \cdots \pi_2$, between the separation pair $(\pi_1, \pi_2) \subset C(e)$ where $\pi_1, \pi_2 \in S$.*

3.2.1 Finding an articulation-flip.

Let (x, y) be any edge inducing a cycle C in $T \cup \{(x, y)\}$ that separates f_u from f_v, let $\pi(u) = \mathrm{m}_T(u, x, y)$ be the projection of u on C.

Now the articulation-flip cases are not necessarily symmetrical. First we present how to detect an articulation-flip, given u, v, and f_u, if f_v plays the role of f_2 (see Section 2.5.1).

If the flip-component containing v is an articulation-flip component, then $\pi(v)$ is an articulation point incident to both f_u and f_v, but the opposite is not necessarily the case. Assume $\pi(v)$ is incident to both f_u and f_v and let c_u denote a corner between $\pi(v)$ and f_u.

Note that if $\pi(v)$ is an articulation point with corners c_1, c_2 both incident to f_v, then f_v is an articulation point in the dual graph with corners c_1, c_2 both incident to $\pi(v)$. Removing f_v from the dual graph would split its

component into several components, and clearly, aside from f_v, only faces in *one* of these components may contain faces incident to v. Any path in the co-tree starting and ending in different components w.r.t the split will have the property that the first face incident to v on that path is f_v. (See Figure 22.)

Now, in the case $f_u = f_1$ and $f_v = f_2$, to find the corner of $\pi(v)$ incident to f_u, we can simply use linkable$(\pi(v), u)$ from before, which will return a corner of f_u incident to $\pi(v)$. To find the two corners of f_v: With the dual structure (see Observation 5) we may mark the face f_v, and expose the vertices u, v. Now, $\pi(v)$ has a unique place in the face-list of some cluster — if and only if that place is in the root cluster, and $n_{\min} = s_{\min} = 0$ for that cluster, f_v plays the role of f_2. That is, if and only if $\pi(v)$ has a corner incident to f_v to one side, and a corner incident to f_v to the other side. In affirmative case, $\pi(v)$ appears with at least one corner to either vertex list; those corners can now be used as cutting-corners for the articulation-flip.

If instead f_v played the role of f_1, a similar procedure is done with $\pi(u)$.

Theorem 4. *Given u, v are not already linkable, we can determine whether u, v are linkable via an articulation-flip in time $\mathcal{O}(\log^2 n)$.*

3.2.2 Finding a separation-flip

Assume v, u are not linkable via an articulation-flip, determine if they are linkable via a separation-flip.

Lemma 16. *Let (x, y) be any edge inducing a cycle C in $T \cup \{(x, y)\}$ that separates f_u from f_v, let $\pi(u) = \mathrm{m}_T(u, x, y)$ be the projection of u on C. Let e_1, e_2 be the edges incident to $\pi(u)$ on C. Then at least one of e_1, e_2 is in the same flip-component as u w.r.t f_u and f_v.*

Proof. This follows from Observation 8: If $\pi(u)$ is incident to both f_u and f_v, then exactly one of the edges e_1, e_2 is in the same flip-component as u. Otherwise, both of the edges e_1, e_2 are in the same flip-component as u. □ □

Lemma 17. *Let C be any fundamental cycle separating f_u from f_v, let e_u be an edge on C in the same flip-component as u, let f_1 be the face adjacent to e_u that is separated from f_u by C, and let $f_2 \in \{v_L, v_R\}$ be a face on the same side of C as f_1. Then f_v is the first face on $f_1 \cdots f_2$ that contains v.*

Proof. $C \cup E[f_v]$ separates f_1 and f_2, so by Observation 6 f_v is on the $f_1 \cdots f_2$ path. It must be the first face on that path that contains v because for any

face f after that, $C \cup V[f]$ does not separate f_1 and f_2, since it can only touch the part of C between $u' = \mathrm{m}(u, x, y)$ and $v' = \mathrm{m}(v, x, y)$ where (x, y) is the edge inducing C. □ □

3.3 Finding the separation pair and corners

Assume u, v are not linkable and not linkable via an articulation-flip.

Lemma 18. *Given u, v, f_u, and f_v, let $(x, y)^*$ be any edge on $f_u \cdots f_v$ inducing a separating cycle C. If $\pi(u) = \pi(v) = \alpha$, then α is one of the separation points if it is adjacent to both f_u and f_v, and otherwise no separation pair for u, v exists. The other separation point, β, is then the first vertex $\neq \alpha$ adjacent to both f_u and f_v on either $\alpha \cdots x$ or $\alpha \cdots y$. If instead $\pi(u) \neq \pi(v)$, then α, β are among the first two vertices adjacent to both f_u and f_v either on $\pi(u) \cdots x$ and $v \cdots x$, or on $u \cdots y$ and $v \cdots y$.*

Proof. If the projection of u equals the projection of v, but u and v are in different flip-components, then the next point incident to both f_u and f_v along the cycle to either side will be the one we are looking for. However, (x, y) may be internal in the flip component containing u or that containing v, and thus one of the searches may return the empty list. But then the other will return the desired pair of vertices.

If the projections are different, and do not themselves form the desired pair (α, β), then we may assume without loss of generality that $\pi(v)$ does not belong to the flip-component containing u. Let X_v, X_u denote the flip-components containing u and v, respectively. If (x, y) is in X_v, such that no edge on $\pi(v) \cdots x$ is incident to both f_u and f_v, then the first vertex on $\pi(v) \cdots y$ incident to f_u and f_v is α. Recall (Observation 8) that $\pi(X_u)$ is an arc $\pi_1 \cdots \pi_2 \subset C$, and suppose without loss of generality π_1 is on the path $u \cdots v$. If $\pi(u) = \pi_1$, β is the second vertex on the path u to y incident to both f_u and f_v, as $\pi(u)$ itself is the first. Otherwise, the first such vertex on the path is β. If, on the other hand, (x, y) did not belong to X_v, let x be the vertex of x, y with the property that the path $u \cdots x$ goes through $\pi(u)$. Then the first vertices on the paths to x which are incident to f_u and f_v both, will be the desired separation pair. □ □

Lemma 19. *In the scenario above, we may find the first two vertices on the path incident to both faces in time $\mathcal{O}(\log^2 n)$.*

Proof. We use the dual structure (see Observation 5) to search for vertices incident to f_u and f_v. Now since the path $\pi(u) \cdots x$ is a sub-path of the fundamental cycle C induced by (x, y) which separates f_u from f_v, all corners

incident to f_v will be on one side, and all corners incident to f_u will be on the other side of the path, or at the endpoints. Thus, we expose f_u and f_v in the dual structure, which takes time $\mathcal{O}(\log^2 n)$. Now expose $\pi(u), x$ in the primal tree. Since this path is part of the separating cycle, if $n_{\min} = s_{\min} = 0$, then the maintained vertex-list will contain exactly those vertices incident to both faces, and a corner list for each of them. We now deal separately with the endpoints exactly as with linkable, by exposing the endpoint faces one by one in the dual structure, and noting whether $c_{\min} = 0$ and in that case, the corner list, for each endpoint. □ □

We conclude with the following theorem.

Theorem 5. *We can maintain an embedding of a dynamic graph under* insert *(edge),* remove *(edge),* cut *(vertex),* join *(vertex), and* flip *(subgraph), together with queries that*

1. (linkable) *Answer whether an edge can be inserted between given endpoints with no other changes to the embedding, and if so, where.*

2. (one-flip-linkable) *Answer whether there exists a flip that would change the answer for* linkable *from "no" to "yes", and if so, what flip.*

The worst case time per operation is $\mathcal{O}(\log^2 n)$.

Acknowledgments

We would like to thank Christian Wulff-Nilsen and Mikkel Thorup for many helpful and interesting discussions and ideas.

References

[1] Stephen Alstrup, Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Transactions on Algorithms*, 1(2):243–264, October 2005.

[2] Giuseppe Di Battista and Roberto Tamassia. Incremental planarity testing. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 436–441. IEEE.

[3] Giuseppe Di Battista and Roberto Tamassia. On-line planarity testing. *SIAM Journal on Computing*, 25:956–997, 1996.

[4] David Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, pages 599–608, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

[5] David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Thomas H. Spencer. Separator based sparsification: I. planarity testing and minimum spanning trees. *Journal of Computer and Systems Sciences*, 52(1):3–27, February 1996.

[6] David Eppstein, Giuseppe F Italiano, Roberto Tamassia, Robert E Tarjan, Jeffery Westbrook, and Moti Yung. Maintenance of a minimum spanning forest in a dynamic plane graph. *Journal of Algorithms*, 13(1):33 – 54, 1992.

[7] Greg N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM Journal on Computing*, 14(4):781–798, 1985.

[8] Zvi Galil, Giuseppe F. Italiano, and Neil Sarnak. Fully dynamic planarity testing with applications. *Journal of the ACM*, 46:28–91, 1999.

[9] John Hopcroft and Robert E. Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, October 1974.

[10] Giuseppe F. Italiano, Johannes A. La Poutré, and Monika H. Rauch. Fully dynamic planarity testing in planar embedded graphs. In Thomas Lengauer, editor, *European Symposium on Algorithms—ESA '93: First Annual European Symposium Bad Honnef, Germany September 30–October 2, 1993 Proceedings*, pages 212–223, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.

[11] David R. Karger. Random sampling in cut, flow, and network design problems. *Mathematics of Operations Research*, pages 648–657, 1994.

[12] Philip N. Klein. Multiple-source shortest paths in planar graphs. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 146–155, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[13] Johannes A. La Poutré. Alpha-algorithms for incremental planarity testing (preliminary version). In *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing*, STOC '94, pages 706–715, New York, NY, USA, 1994. ACM.

[14] Mihai Pătraşcu and Erik D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM Journal on Computing*, 35(4):932–963, 2006. See also STOC'04, SODA'04.

[15] Mihai Patrascu and Mikkel Thorup. Planning for fast connectivity updates. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 263–271, Washington, DC, USA, 2007. IEEE Computer Society.

[16] Daniel D. Sleator and Robert E. Tarjan. A data structure for dynamic trees. *Journal of Computer and Systems Sciences*, 26(3):362–391, June 1983.

[17] William T. Tutte. How to Draw a Graph. *Proceedings of the London Mathematical Society*, s3-13(1):743–767, 1963.

[18] William T. Tutte. *Graph theory.* Encyclopedia of mathematics and its applications. Addison-Wesley Pub. Co., Advanced Book Program, 1984.

[19] Jeffery Westbrook. Fast incremental planarity testing. In W. Kuich, editor, *Automata, Languages and Programming (ICALP)*, volume 623 of *Lecture Notes in Computer Science*, pages 342–353. Springer Berlin Heidelberg, 1992.